

What Is Claimed

1. A method of asynchronous data replication over a network, wherein one or more applications send write commands to a primary host, comprising:
at the primary host:

- acknowledging each write command before transmitting the write command to a secondary host;
- assigning a data set ID to each write command;
- transmitting each write command before the application has sent all of the write commands of the data set to the primary host;
- assigning the data set ID to a sync command;
- transmitting the sync command, wherein the write commands and the sync command define a data set;

at the secondary host:

- receiving the write commands and sync command out of order; and
- writing the data set to a storage coupled to the secondary host.

2. The method of claim 1, further comprising tracking the receipt of write and sync commands at the secondary host.

3. The method of claim 2, further comprising assigning a unique sequence number for each write command and another sequence number to the sync command at the primary host.

4. The method of claim 3, wherein tracking the receipt of write and sync commands includes setting bits in an identifier bitmap corresponding to the data set at the secondary host.

5. The method of claim 1, further comprising writing the blocks to a secondary target VLUN and acknowledging the write and sync commands before receipt of the data set at secondary host.

6. The method of claim 1, further comprising storing multiple versions of one of the blocks at the primary host while waiting for an acknowledgement from the secondary host.

7. The method of claim 4, further comprising using the identifier bitmap to determine receipt of the data set at the secondary host.

8. The method of claim 3, further comprising updating a data structure at the secondary host to track the blocks in the secondary target VLUN.

9. The method of claim 8, wherein updating the data structure at the secondary host includes setting bits in a bitmap and updating a log file.

10. The method of claim 1, further comprising writing the blocks to a memory buffer at the secondary host and acknowledging the write and sync commands at the secondary host before receipt of the data set.

11. The method of claim 1, wherein the step of receiving at the secondary host includes receiving write commands belonging to a plurality of data sets that are overlapping with respect to each other.

12. The method of claim 1, further comprising receiving a second data set at the secondary host, wherein the step of receiving write commands for the second data set precedes receiving write commands of the data set.

13. The method of claim 1, further comprising storing the data sets in a plurality of data structures.

14. The method of claim 1, further comprising transmitting the write and sync commands on the network using the Fibre channel protocol.

15. The method of claim 1, further comprising transmitting the write and sync commands on the network using the iSCSI protocol.

16. The method of claim 1, wherein the step of writing includes the secondary host writing the blocks of the data set to the secondary source VLUN.

17. The method of claim 1, wherein the data set contains one or more file system transactions.

18. A system of asynchronous data replication, comprising:
a primary host asynchronously transmitting write commands and a sync command, wherein each write command has a data set ID and each sync command has a data set ID; and

a secondary host, coupled to the primary host, tracking receipt of write and sync commands and after receipt of the write and sync commands of the data set ID writing the write commands to a storage coupled to the secondary host.

19. The system of claim 18, further comprising a data structure at the secondary host for tracking receipt of write and sync commands.

20. The system of claim 19, wherein the data structure includes a flag to indicate when the write commands and the sync command of a data set are received at the secondary host.

21. The system of claim 20, wherein the secondary host writes the data set to storage after a temporary failure of the secondary host if the flag indicates that the write and sync commands of the data set were received.

22. A system of asynchronous data replication, comprising:
a primary host asynchronously transmitting write and sync commands for a first data set and a second data set; and
a secondary host, coupled to the primary host, receiving the write and sync commands and writing the data set to storage, wherein the primary host transmits write and sync commands for the second data set before the secondary host acknowledges the first data set was written to storage.

23. A system of asynchronous data replication, comprising:
a primary host asynchronously transmitting write commands and the blocks without waiting to complete a data set followed by a sync command; and
a secondary host coupled to the network receiving the write commands, the blocks, and the sync command in any order, wherein the secondary host writes the data set to storage coupled to the secondary host.

24. A system of asynchronous data replication, comprising:
a primary host receiving write commands from an application, generating sync commands, and transmitting the write and sync commands over a network asynchronously and independently from one or more data sets; and
a secondary host, coupled to the network, receiving the write and sync commands out of order and writing the data sets in order to storage coupled to the secondary host.

25. A method of data replication, comprising the steps of:
transmitting a data set from a primary host to a secondary host wherein the secondary host includes a data structure with a complete flag;
writing the data set to a target VLUN at the secondary host;
setting the complete flag = yes once the data set is written to the target VLUN at the secondary host; and
writing the data set to a source VLUN at the secondary host.

26. A method of data replication comprising the steps of:
transmitting a data set from a primary host to a secondary host wherein the secondary host includes a data structure with a complete flag;
writing the data set to a target VLUN at the secondary host;
setting the complete flag = yes once the data set is written to the target VLUN at the secondary host;
resuming operation of the secondary host after a failure; and
writing the data set to a source VLUN at the secondary host.

27. A method of data replication comprising the steps of:
transmitting a data set from a primary host to a secondary host wherein the secondary host includes a data structure with a complete flag;
writing part of the data set to a target VLUN at the secondary host;
resuming operation of the secondary host after a failure;
erasing the part of data set written to the target VLUN at the secondary host;
retransmitting the data set from the primary host to the secondary host;
writing the data set to the target VLUN at the secondary host;
setting the complete flag = yes once the data set is written to the target VLUN at the secondary host; and
writing the data set to a source VLUN at the secondary host.